

WHAT IS CLAIMED IS:

1 1. A method of implementing an in-band signaling
2 mechanism in a resilient packet ring ("RPR") network
3 comprising a plurality of nodes and first and second
4 counter-rotating rings each comprising a plurality of
5 links for carrying information between the nodes, the
6 method comprising the step of:

7 transmitting a data packet on the first ring of the
8 RPR network, wherein the data packet includes at least
9 one field requesting information on traffic flow from a
10 downstream node; and

11 receiving the data packet on the second ring wherein
12 the at least one field indicates the status of traffic
13 flow in response to the request for information.

1 2. A method of implementing an in-band signaling
2 mechanism in a resilient packet ring ("RPR") network
3 comprising a plurality of nodes and first and second
4 counter-rotating rings each comprising a plurality of
5 links for carrying information between the nodes, the
6 method comprising the step of:

7 including in each packet transmitted on the RPR
8 network a header comprising at least one fairness
9 specific field for use in detecting and correcting lack
10 of fairness among the nodes, at least one load balancing
11 specific field for use in detecting and correcting load
12 imbalances between the rings, at least one congestion
13 control specific field for use in detecting and
14 alleviating congestion on a ring, and at least one common
15 field for use in detecting and correcting lack of
16 fairness, load imbalance, and congestion in the RPR.

1 3. The method of claim 2 wherein the at least one
2 fairness specific field is selected from the group
3 consisting of a field for indicating that a lack of
4 fairness has been detected between two nodes and a field
5 for indicating a number of flows belonging to a specified
6 class of service at a node.

1 4. The method of claim 2 wherein the at least one
2 load balancing specific field is selected from a group
3 consisting of a field for indicating that the packet is
4 a test packet, a field for indicating that the packet
5 must be inspected for load balancing purposes, and a
6 field for indicating failure on one or more links
7 connected to a node.

1 5. The method of claim 2 wherein the at least one
2 congestion control specific field is selected from a
3 group consisting of a field for indicating that
4 congestion has been detected with respect to a given
5 class of service and a field for indicating a frequency
6 with which packets are to be marked by a node for
7 congestion control inspection purposes.

1 6. The method of claim 2 wherein the at least one
2 common field is selected from a group consisting of a
3 field for indicating a class of service of the packet, a
4 field for indicating available bandwidth for the class of
5 service of the packet, a field for indicating an
6 allocated bandwidth for the class of service of the
7 packet, a field for indicating a time at which the packet
8 left a first node destined for a second node, a field for
9 indicating a time at which the packet was received at the
10 second node, a field for indicating that the packet is to
11 be time stamped for the class of service; and a field for
12 indicating that the packet has been received on a counter
13 ring.

1 7. The method of claim 2 wherein the RPR network
2 is a wavelength division multiplex RPR ("WDMRPR").

1 8. The method of claim 7 wherein the first and
2 second rings comprise first and second wavelengths,
3 respectively.

1 9. A method of implementing an in-band signaling
2 mechanism for detecting and correcting congestion on a
3 first ring of a resilient packet ring ("RPR") network,
4 the method comprising the steps of, for each nth packet
5 of a flow of a first class of service at a first node:

6 time stamping the packet with a first time
7 indicative of a time the packet is sent to a second node;

8 sending the packet to the second node on a first
9 link between the first and second nodes, the first link
10 comprising a portion of the first ring;

11 upon receipt by the second node of the packet, time
12 stamping the packet with a second time indicative of a
13 time the packet was received at the second node;

14 calculating a difference between the first and
15 second times; and

16 responsive to the difference being greater than a
17 predetermined threshold, signaling to the first node that
18 congestion has been detected on the first link between
19 the first and second nodes.

1 10. The method of claim 9 further comprising, prior
2 to the step of time stamping the packet with a first time
3 indicative of a time the packet is sent to a second node,
4 the step of setting a designated bit in a header of the
5 packet.

1 11. The method of claim 9 wherein the step of time
2 stamping is performed by indicating a time in a
3 designated field of a header of the packet.

1 12. The method of claim 9 further comprising the
2 steps of:

3 responsive to the difference being greater than the
4 predetermined threshold, signaling a BB to reduce the
5 available bandwidth for the first class of service to
6 reduce the number of new flows admitted on the first link
7 between the first and second nodes; and

8 responsive to the signaling, reducing the token
9 bucket rate at the first node for the first class of
10 service.

1 13. The method of claim 9 further comprising the
2 step of, responsive to the difference being greater than
3 the predetermined threshold, reducing n at the first
4 node, thereby increasing the number of packets of traffic
5 of the first class of service that are time stamped with
6 a first time.

1 14. The method of claim 9 further comprising the
2 step of reducing the token bucket rate at the second node
3 for the first class of service responsive to detection of
4 congestion on the first link.

1 15. The method of claim 9 further comprising the
2 steps of, responsive to the difference being less than
3 the predetermined threshold, signaling a BB to increase
4 the available bandwidth for the first class of service to
5 increase the number of new flows admitted on the first
6 link between the first and second nodes, and increasing
7 the token bucket rate at the first node for the first
8 class of service.

1 16. The method of claim 9 wherein the step of
2 signaling comprises the steps of:

3 setting a bit in a header of the packet to indicate
4 to the first node that congestion has been detected on
5 the first link; and

6 sending the packet back to the first node on a
7 second link between the first and second nodes, the
8 second link comprising a portion of a second ring of the
9 RPR, wherein the second link is a reverse link of the
10 first link.

1 17. The method of claim 9 wherein the RPR network
2 is a wavelength division multiplex RPR ("WDMRPR").

1 18. The method of claim 17 wherein the first and
2 second rings comprise first and second wavelengths,
3 respectively.
4

1 19. A method of implementing an in-band signaling
2 mechanism for detecting and correcting load imbalance
3 between first and second counter-rotating rings of a
4 resilient packet ring ("RPR") network comprising a
5 plurality of nodes, the method comprising the steps of:

6 time stamping a packet at a first node with a first
7 time value indicative of a time the packet is sent from
8 the first node to a second node via the first ring;

9 sending the packet to the second node via the first
10 ring;

11 at the second node, time stamping the packet with a
12 second time value indicative of a time the packet is
13 received by the second node via the first ring;

14 returning the packet to the first node via the
15 second ring;

16 at the first node, recording a difference between
17 the first and second time values;

18 time stamping the packet at the first node with a
19 third time value indicative of a time the packet is sent
20 from the first node to the second node via the second
21 ring;

22 sending the packet to the second node via the second
23 ring;

24 at the second node, time stamping the packet with a
25 fourth time value indicative of a time the packet is
26 received by the second node via the second ring;

27 returning the packet to the first node via the first
28 ring;

29 at the first node, recording a difference between
30 the third and fourth time values; and
31 responsive to a determination that an absolute value
32 of the difference between the third and fourth time
33 values is less than an absolute value of the difference
34 between the first and second time values, signaling a
35 bandwidth broker ("BB") perform load balancing between
36 the first and second rings.

1 20. The method of claim 19 further comprising,
2 prior to the step of sending the packet to the second
3 node via the first ring, the step of setting a designated
4 bit in a header of the packet for indicating to the
5 second node that the packet is to be examined for load
6 balancing purposes.

1 21. The method of claim 19 further comprising,
2 prior to the step of returning the packet to the first
3 node via the second ring, the step of setting a
4 designated bit in a header of the packet for indicating
5 to the first node that the packet was sent via a counter
6 ring.

1 22. The method of claim 19 further comprising,
2 prior to the step of sending the packet to the second
3 node via the second ring, the step of setting a
4 designated bit in a header of the packet for indicating
5 to the second node that the packet is a test packet.

1 23. The method of claim 19 wherein the step of time
2 stamping is performed by indicating a time in a
3 designated field of a header of the packet.

1 24. The method of claim 19 wherein the step of
2 signaling the BB to perform load balancing comprises the
3 step of signaling the BB to reroute some of the existing
4 traffic via the second ring.

1 25. The method of claim 19 wherein the step of
2 signaling the BB to perform load balancing comprises the
3 step of signaling the BB to admit new incoming flows to
4 the first node on the second ring.

1 26. The method of claim 19 wherein load balancing
2 is performed at periodic intervals.

1 27. The method of claim 19 wherein load balancing
2 is performed responsive to failure of one or more of the
3 rings.

1 28. The method of claim 19 wherein load balancing
2 is performed when a new flow enters the RPR network.

1 29. The method of claim 28 wherein failure of one
2 or more of the rings is indicated in a designated field
3 of a header of the packet.

1 30. The method of claim 19 wherein the RPR network
2 is a wavelength division multiplex RPR ("WDMRPR").

1 31. The method of claim 30 wherein the first and
2 second rings comprise first and second wavelengths,
3 respectively.

1 32. A method of using an in-band signaling
2 mechanism for detecting and correcting a lack of fairness
3 between an upstream node and a downstream node with
4 respect to traffic of a first class of service in a
5 resilient packet ring ("RPR") network comprising first
6 and second two counter-rotating rings, the method
7 comprising the steps of:

8 causing a bandwidth broker ("BB") to increase the
9 allocated bandwidth and number of admissible flows for
10 the first class of service at the downstream node; and

11 signaling the upstream node to reduce the amount of
12 traffic of the first class of service being sent from the
13 upstream node to the downstream node.

1 33. The method of claim 32 wherein the step of
2 signaling the upstream node further comprises the step of
3 signaling the BB to decrease a flow rate for each flow of
4 the first class of service at the upstream node to a
5 predetermined value.

1 34. The method of claim 32 wherein the step of
2 signaling the upstream node further comprises the step of
3 signaling to the BB to admit only new flows for the first
4 class of service having a peak rate of less than a
5 selected value.

1 35. The method of claim 32 wherein the step of
2 signaling the upstream node comprises the step of setting
3 a designated bit in a header of the packet.

1 36. The method of claim 32 wherein the RPR network
2 is a wavelength division multiplex RPR ("WDMRPR").

1 37. The method of claim 36 wherein the first and
2 second rings comprise first and second wavelengths,
3 respectively.

1 38. An in-band signaling mechanism for detecting
2 and correcting congestion on a first ring of a resilient
3 packet ring ("RPR") network, the mechanism comprising:

4 means for time stamping each nth packet of a flow of
5 a first class of service at a first node with a first
6 time indicative of a time the packet is sent to a second
7 node;

8 means for sending the packet to the second node on
9 a first link between the first and second nodes, the
10 first link comprising a portion of the first ring;

11 means for time stamping the packet with a second
12 time indicative of a time the packet was received at the
13 second node upon receipt by the second node of the
14 packet;

15 means for calculating a difference between the first
16 and second times; and

17 means responsive to the difference being greater
18 than a predetermined threshold for signaling to the first
19 node that congestion has been detected on the first link
20 between the first and second nodes.

1 39. The mechanism of claim 38 wherein the RPR
2 network is a wavelength division multiplex RPR
3 ("WDMRPR").

1 40. The mechanism of claim 39 wherein the first
2 ring comprises a first wavelength.

1 41. An in-band signaling mechanism for detecting
2 and correcting load imbalance between first and second
3 counter-rotating rings of a resilient packet ring ("RPR")
4 network comprising a plurality of nodes, the mechanism
5 comprising:

6 means for time stamping a packet at a first node
7 with a first time value indicative of a time the packet
8 is sent from the first node to a second node via the
9 first ring;

10 means for sending the packet to the second node via
11 the first ring;

12 means for time stamping the packet with a second
13 time value indicative of a time the packet is received by
14 the second node via the first ring;

15 means for returning the packet to the first node via
16 the second ring;

17 means for recording a difference between the first
18 and second time values;

19 means for time stamping the packet at the first node
20 with a third time value indicative of a time the packet
21 is sent from the first node to the second node via the
22 second ring;

23 means for sending the packet to the second node via
24 the second ring;

25 means for time stamping the packet with a fourth
26 time value indicative of a time the packet is received by
27 the second node via the second ring;

28 means for returning the packet to the first node via
29 the first ring;

30 means for recording a difference between the third
31 and fourth time values; and

32 means responsive to a determination that an absolute
33 value of the difference between the third and fourth time
34 values is less than an absolute value of the difference
35 between the first and second time values for signaling a
36 bandwidth broker ("BB") perform load balancing between
37 the first and second rings.

1 42. The mechanism of claim 41 wherein the RPR
2 network is a wavelength division multiplex RPR
3 ("WDMRPR").

1 43. The mechanism of claim 42 wherein the first and
2 second rings comprise first and second wavelengths,
3 respectively.

1 44. An in-band signaling mechanism for detecting
2 and correcting a lack of fairness between an upstream
3 node and a downstream node with respect to traffic of a
4 first class of service in a resilient packet ring ("RPR")
5 network comprising first and second two counter-rotating
6 rings, the mechanism comprising:

7 means for causing a bandwidth broker ("BB") to
8 increase the allocated bandwidth and number of admissible
9 flows for the first class of service at the downstream
10 node; and

11 means for signaling the upstream node to reduce the
12 amount of traffic of the first class of service being
13 sent from the upstream node to the downstream node.

1 45. The mechanism of claim 44 wherein the RPR
2 network is a wavelength division multiplex RPR
3 ("WDMRPR").

1 46. The mechanism of claim 45 wherein the first and
2 second rings comprise first and second wavelengths,
3 respectively.